

MACHINE LEARNING METHOD

Abstract

A method for using machine learning to solve problems having either a "positive" result (the event occurred) or a "negative" result (the event did not occur), in which the probability of a positive result is very low and the consequences of the positive result are significant. Training data is obtained and a subset of that data is distilled for application to a machine learning system. The training data includes some records corresponding to the positive result, some nearest neighbors from the records corresponding to the negative result, and some other records corresponding to the negative result. The machine learning system uses a co-evolution approach to obtain a rule set for predicting results after a number of cycles. The machine system uses a fitness function derived for use with the type of problem, such as a fitness function based on the sensitivity and positive predictive value of the rules. The rules are validated using the entire set of training data.

15

20

25

30